arXiv:1703.07809v1 [math.NA] 22 Mar 2017

Empirical Risk Minimization as Parameter Choice Rule for General Linear Regularization Methods

Housen Li
housen.li@outlook.com
National University of Defense Technology, Changsha, China
Frank Werner[1]
frank.werner@mpibpc.mpg.de
Max Planck Institute for Biophysical Chemistry, Göttingen, Germany
and
Felix Bernstein Institute for Mathematical Statistics in the Bioscience,
University of Goettingen, Germany

We consider the statistical inverse problem to recover f from noisy measurements $Y = Tf + \sigma\xi$ where ξ is Gaussian white noise and T a compact operator between Hilbert spaces. Considering general reconstruction methods of the form $\hat{f}_\alpha = q_\alpha(T^*T)T^*Y$ with an ordered filter q_α, we investigate the choice of the regularization parameter α by minimizing an unbiased estimate of the predictive risk $\mathbb{E}\left[\|Tf - T\hat{f}_\alpha\|^2\right]$. The corresponding parameter α_{pred} and its usage are well-known in the literature, but oracle inequalities and optimality results in this general setting are unknown. We prove a (generalized) oracle inequality, which relates the direct risk $\mathbb{E}\left[\|f - \hat{f}_{\alpha_{\text{pred}}}\|^2\right]$ with the oracle prediction risk $\inf_{\alpha>0}\mathbb{E}\left[\|Tf - T\hat{f}_\alpha\|^2\right]$. From this oracle inequality we are then able to conclude that the investigated parameter choice rule is of optimal order.

Finally we also present numerical simulations, which support the order optimality of the method and the quality of the parameter choice in finite sample situations.

Keywords: statistical inverse problem, regularization method, filter-based inversion, a-posteriori parameter choice rule, order optimality, exponential bounds, oracle inequality.

AMS classification numbers: Primary 62G05; Secondary 62G20, 65J22, 65J20.

1 Introduction

Suppose we want to recover an unknown function $f \in \mathcal{X}$ from noisy measurements

$$Y = Tf + \sigma\xi \tag{1}$$

where $T : \mathcal{X} \to \mathcal{Y}$ is an operator between Hilbert spaces $\mathcal{X}$ and $\mathcal{Y}$, ξ is a standard Gaussian white noise process and $\sigma > 0$ denotes the noise level. In fact, (1) has to be understood in a weak sense as $\xi \notin \mathcal{Y}$, i.e. for each $y \in \mathcal{Y}$ we have access to observations of the form

$$Y_y := \langle Tf, y\rangle_{\mathcal{Y}} + \sigma\langle \xi, y\rangle_{\mathcal{Y}^*\times\mathcal{Y}}$$

where $\langle \xi, y\rangle \sim \mathcal{N}\left(0, \|y\|_{\mathcal{Y}}^2\right)$ and $\mathbb{E}\left[\langle \xi, y_1\rangle\langle \xi, y_2\rangle\right] = \langle y_1, y_2\rangle$ for all $y_1, y_2 \in \mathcal{Y}$. Models of the form (1) underly a plenitude of applications, see e.g. (O'Sullivan, 1986), and have been considered by Mathé and Pereverzev (2001); Bissantz et al. (2007); Ingster et al. (2012, 2014); Werner (2015).

[1] Corresponding author

Throughout the paper we will assume that the operator T is injective, compact and Hilbert-Schmidt, i.e. the squares of its singular values are summable. Especially, this implies that its singular values tend to 0 and hence the inversion of T is unstable, i.e. the problem to recover f from (1) is ill-posed and regularization is needed, see Cavalier (2011); Engl et al. (1996) and the references therein. In the literature many different approaches for estimation of f can be found, including methods based on a singular value decomposition (SVD) of T, see e.g. (Johnstone and Silverman, 1991; Mair and Ruymgaart, 1996; Johnstone et al., 2004; Cavalier and Golubev, 2006), wavelet-vaguelette (Donoho, 1995) and vaguelette-wavelet methods (Abramovich and Silverman, 1998), and Galerkin-type methods (Cohen et al., 2004).

In this paper we follow a common paradigm and consider regularization methods written in terms of an *ordered filter* $q_\alpha : [0, \|T^*T\|] \to \mathbb{R}$ parametrized by $\alpha > 0$ (see e.g. Definition 1 for the precise requirements on q_α), meaning that the regularized solution is given by

$$\hat{f}_\alpha = q_\alpha\left(T^*T\right) T^*Y. \tag{2}$$

Regularization methods of the form (2) include famous examples like spectral cut-off and Tikhonov regularization, and have been studied extensively in the literature, see (Engl et al., 1996) and the references therein, and (Bissantz et al., 2007).

Choosing an appropriate parameter α in (2) is an important problem in regularization theory as it dramatically influences the performance of the estimator $\hat{f}_\alpha$. Bissantz et al. (2007) show that estimators of the form (2) are order-optimal over certain smoothness classes $\mathcal{W} \subset \mathcal{X}$, if the parameter $\alpha = \alpha^*$ is chosen in a reasonable a-priori way (depending on properties of $\mathcal{W}$). This means $\hat{f}_{\alpha^*}$ achieves the best possible rate of convergence w.r.t. the *direct risk* $R(\alpha, f) := \mathbb{E}\left[\left\|\hat{f}_\alpha - f\right\|_{\mathcal{X}}^2\right]$ in these classes. In practice, the parameter α has to be chosen without any knowledge of f (and hence of $\mathcal{W}$), which makes a-priori parameter choice rules useless. Therefore, *a-posteriori* parameter choice rules are of interest, as they make only use of the data Y and the noise level $\sigma > 0$. As a-posteriori parameter choice rules have to adapt to the unknown properties of $\mathcal{W}$ (and hence of f), this issue is also known as *adaptivity*. For simplicity, we will assume here that σ is known. In practice, the parameter σ can typically be estimated sufficiently fast from the data or the measurement setting, see e.g. (Rice, 1984; Hall et al., 1990; Dette et al., 1998). We will discuss this situation in more detail in Section 6.

A variety of a-posteriori parameter choice rules have been proposed in the literature, including the discrepancy principle (Morozov, 1966; Davies and Anderssen, 1986; Lukas, 1995; Blanchard et al., 2016), generalized cross-validation (Wahba, 1977; Golub et al., 1979; Lukas, 1993), the Lepskiĭ-type balancing principle (Lepskiĭ, 1991; Mathé, 2006; Mathé and Pereverzev, 2006; Werner and Hohage, 2012) and many more. We refer to (Bauer and Lukas, 2011) for a recent overview and numerical comparison. General adaptivity in statistical inverse problems has also been treated in (Goldenshluger, 1999; Tsybakov, 2000; Cavalier et al., 2003).

In this paper we deal with a specific method based on empirical risk minimization originally introduced by Mallows (1973) for model selection in linear regression and therefore known as *Mallow's C_L*. Consider the *prediction risk* $r(\alpha, f) := \mathbb{E}\left[\left\|T\left(\hat{f}_\alpha - f\right)\right\|_{\mathcal{Y}}^2\right]$. Following Stein (1981), we find that an (up to a constant independent of α) unbiased estimator for this quantity is given by

$$\hat{r}(\alpha, Y) := \left\|T\hat{f}_\alpha\right\|_{\mathcal{Y}}^2 - 2\langle T\hat{f}_\alpha, Y\rangle + 2\sigma^2 \mathrm{Trace}\left(T^*T q_\alpha\left(T^*T\right)\right). \tag{3}$$

Now the idea is to choose α as a minimizer of $\hat{r}(\alpha, Y)$, i.e.

$$\alpha_{\mathrm{pred}} \in \operatorname*{argmin}_{\alpha>0} \hat{r}(\alpha, Y).$$

Note that the functional (3) penalizes the the misfit between the model $\hat{f}_\alpha$ and the data Y, and furthermore $2\sigma^2 \mathrm{Trace}\left(T^*T q_\alpha\left(T^*T\right)\right)$ penalizes the number of degrees of freedom of the model.

We refer to Birgé and Massart (2001, 2007) for details and a discussion of other possible penalty terms in (3).

It is known that choosing $\alpha = \alpha_{\text{pred}}$ in combination with certain regularization schemes leads to an order optimal method w.r.t. the prediction risk $r(\alpha, f)$, see e.g. (Li, 1987; Vogel, 1986; Lukas, 1993). A very precise result about its performance which is also a central ingredient of this paper can be found in the seminal paper by Kneip (1994), who proves exponential deviation bounds for $r(\alpha_{\text{pred}}, f)$.

Due to ill-posedness, optimality w.r.t. the prediction risk is however a very weak statement, and consequently we are interested in order optimality w.r.t. the direct risk. This question has hardly been touched in literature, and to the authors best knowledge the only result is due to Chernousova and Golubev (2014) who restrict to finite dimensional spaces and spectral cut-off regularization. Nevertheless, the choice $\alpha = \alpha_{\text{pred}}$ has successfully been applied in image denoising applications, see e.g. (Luisier et al., 2007; Candès et al., 2013; Wang and Morel, 2013; Deledalle et al., 2014). Furthermore, the distributional behavior of α_{pred} has recently been studied by Lucka et al. (2017). There it has been argued that the choice α_{pred} and choices based on unbiased risk minimization in general do not seem suitable for inverse problems. Besides, it remains one of the most popular parameter selection rules, due to its favorable practical performance (see e.g. Bauer and Lukas (2011); Chernousova and Golubev (2014)). Supporting this fact, we will prove an *oracle inequality* of the form

$$R(\alpha_{\text{pred}}, f) \leq \Theta\left(\min_{\alpha>0} r(\alpha, f)\right) \tag{4}$$

for all $f \in \mathcal{W}$ as $\sigma \searrow 0$ where Θ is some (explicit) functional and $\mathcal{W} \subset \mathcal{X}$ some smoothness class. More specifically, the functional Θ will be such that we can derive order optimality of $\hat{f}_{\alpha_{\text{pred}}}$ under weak assumptions over many classes $\mathcal{W} \subset \mathcal{X}$. Moreover, we provide a general strategy to prove estimates of the form (4) which is of independent interest and might be used for the analysis of other a-posteriori parameter choice rules as well. This makes our analytical methodology substantially different from that in (Chernousova and Golubev, 2014), since the crucial techniques (such as optional stopping of martingales) there do not apply to general regularization schemes. For more details on oracle inequalities in statistical inverse problems we refer to (Cavalier et al., 2002; Blanchard et al., 2016).

The rest of this paper is organized as follows. In the next section we introduce ordered filter based regularization methods and the empirical parameter choice rule via unbiased prediction risk minimization. The convergence analysis of such a rule is established by means of an oracle inequality in Section 3 with corresponding conclusions on order optimality. In Section 4 we present a general methodology for proving oracle inequalities of the form (4) and apply this methodology for proving the results from Section 3. The performance, as well as the convergence behavior, of the analyzed parameter choice rule is examined by comparison with other empirical parameter choice rules in a simulation study in Section 5. We end this paper with some conclusions in Section 6 .

2 Filter based regularization and empirical risk minimization

2.1 Ordered filter based regularization methods

Suppose that $\mathcal{X}$ and $\mathcal{Y}$ are Hilbert spaces, $T : \mathcal{X} \to \mathcal{Y}$ is an injective and compact Hilbert-Schmidt operator, and ξ in (1) is a standard Gaussian white noise process as described in the introduction. To simplify the notation we will always assume that $\dim(\mathcal{X}) = \infty$, but stress that the case of finite dimensional $\mathcal{X}$ (and $\mathcal{Y}$) can be treated similarly. By assumption, there exists a singular value decomposition (SVD) $\left\{\left(\sqrt{\lambda_k}, e_k, g_k\right)\right\}_{k\in\mathbb{N}}$ of T where $\lambda_1 \geq \lambda_2 \geq \cdots > 0$ are the eigenvalues of T^*T and $e_1, e_2, \ldots$ are the corresponding normalized eigenvectors. By introducing the notation $Y_k := \langle g_k, Y\rangle$, $\xi_k := \langle g_k, \xi\rangle$ and $f_k = \langle f, e_k\rangle$, we equivalently transform the model (1) to the

Gaussian sequence model

$$Y_k = \sqrt{\lambda_k} f_k + \sigma \xi_k, \qquad k = 1, 2, \ldots, \tag{5}$$

with $\xi_k \overset{iid}{\sim} \mathcal{N}(0,1)$.

As mentioned in the introduction, we focus on regularization methods of the form (2), which by meanings of (5) can be equivalently formulated as

$$\left(\hat{f}_\alpha\right)_k = \sqrt{\lambda_k} q_\alpha(\lambda_k) Y_k, \qquad k \in \mathbb{N}.$$

Remark 1. *Note that T being Hilbert-Schmidt implies that $\sum_{k=1}^{\infty} \lambda_k < \infty$, and hence T^*Y can be interpreted as a random variable with values in $\mathcal{X}$ as*

$$\mathbb{E}\left[\|T^*\xi\|_{\mathcal{X}}^2\right] = \mathbb{E}\left[\sum_{k=1}^{\infty} \langle \xi, Te_k\rangle_{\mathcal{Y}}^2\right] = \sum_{k=1}^{\infty} \lambda_k \mathbb{E}\left[\langle \xi, g_k\rangle_{\mathcal{Y}}^2\right] < \infty.$$

Consequently, $\hat{f}_\alpha$ as in (2) is well-defined.

Estimators of this form can be understood as stable approximations of the well-known least squares estimate $\hat{f} := (T^*T)^{-1}T^*Y$ (or $\hat{f}_k = Y_k/\sqrt{\lambda_k}$) in the sense of replacing $(\cdot)^{-1}$ with a function $q_\alpha(\cdot)$. To obtain a well-defined and reasonable regularization method, the functions $q_\alpha(\cdot)$ should satisfy proper conditions. We are particularly interested in case that $q_\alpha(\cdot)$ is an ordered filter.

Definition 1. *Let $q_\alpha : [0, \lambda_1] \to \mathbb{R}$, indexed by $\alpha \in \mathcal{A} \subset \mathbb{R}_+$, be a sequence of functions. We always assume that 0 is an accumulation point of $\mathcal{A}$ and that $\mathcal{A}$ is bounded.*

(i) The family $q_\alpha, \alpha \in \mathcal{A}$, is called a filter, if there exist constants $C_q', C_q'' > 0$ such that for every $\alpha \in \mathcal{A}$ and every $\lambda \in [0, \lambda_1]$

$$\alpha\,|q_\alpha(\lambda)| \le C_q' \qquad \text{and} \qquad \lambda\,|q_\alpha(\lambda)| \le C_q''.$$

(ii) The filter $q_\alpha, \alpha \in \mathcal{A}$, is called ordered, if the sequence $\{q_\alpha(\lambda_k)\}_{k=1}^{\infty}$ is strictly monotone, i.e.

$$\alpha < \beta \qquad \Rightarrow \qquad \forall\, k \in \mathbb{N} : q_\alpha(\lambda_k) \le q_\beta(\lambda_k) \text{ and } \exists\, k_0 \in \mathbb{N} \text{ such that } q_\alpha(\lambda_{k_0}) < q_\beta(\lambda_{k_0}),$$

and continuous as $\mathcal{A} \ni \alpha \mapsto \{q_\alpha(\lambda_k)\}_{k=1}^{\infty} \in \ell^2$.

The requirement of an ordered filter is rather weak, as it is satisfied by various regularization methods. In Table 1 we give several examples of such. Note that for spectral cut-off regularization, the set $\mathcal{A}$ has to be chosen as $\{\lambda_k : k = 1, 2, \ldots\}$ in order to guarantee the strict monotonicity and the continuity required by condition (ii) in Definition 1. In Table 1 we also indicate whether the method can be implemented without SVD. This property is extremely crucial in practice, especially for large-scale applications, where the computation of an SVD is often impossible given limited time and resources. The implementation of Showalter's method, for instance, can avoid SVD by employing Runge-Kutta schemes, see e.g. (Rieder, 2005). For a further discussion of these and other methods we refer to the monograph (Engl et al., 1996).

In this paper, we focus on the asymptotic properties of ordered filter based regularization methods as the noise level σ tends to zero. As noticed by Bakushinskiĭ (1984), the convergence rates of regularization methods can be arbitrarily slow if the underlying problem ill-posed. In order to derive rates of convergence we need assume some smoothness about the unknown truth f. Typically, the smoothness of f is measured relative to the smoothing properties of the forward operator T in terms of a *source condition*, i.e. we assume that

$$f \in \mathcal{W}_\phi := \{f \in \mathcal{X} : f = \phi(T^*T)w, \|w\|_{\mathcal{X}} \le C\} \qquad \text{for some constant } C, \tag{6}$$

where ϕ is a so-called *index function*, i.e. ϕ is continuous, strictly increasing, and $\phi(0) = 0$. Note that for any $f \in \mathcal{X}$ there exists a function $\widetilde{\phi}$ such that $f \in \mathcal{W}_{\widetilde{\phi}}$, cf. (Mathé and Hofmann, 2008).

To take advantage of (6) we furthermore assume that ϕ is a *qualification* of the filter q_α, this is

$$\sup_{\lambda\in[0,\lambda_1]} \phi(\lambda)\,|1-\lambda q_\alpha(\lambda)| \le C_\phi \phi(\alpha) \qquad \text{for all } \alpha \in \mathcal{A}, \tag{7}$$

with C_ϕ being a constant depending only on ϕ. For further details on general source conditions and corresponding qualifications we refer to Mathé and Pereverzev (2003).

As an example consider $\phi(t) = t^v$, which is known as Hölder type source condition of order $v > 0$:

$$\mathcal{W}_v := \{f \in \mathcal{X} : f = (T^*T)^v w, \|w\|_{\mathcal{X}} \le C\}\,. \tag{8}$$

The function $\phi(t) = t^v$ is a qualification of the filter q_α if

$$\sup_{\lambda\in[0,\lambda_1]} \lambda^v\,|1-\lambda q_\alpha(\lambda)| \le C_v \alpha^v \qquad \text{for all } \alpha \in \mathcal{A}. \tag{9}$$

In this case, the largest possible v such that (9) is satisfied, is called the classical or polynomial *qualification index* v_0 of the ordered filter q_α. For the methods discussed in Table 1, v_0 as well as C_v is also depicted.

Table 1: Examples of ordered filters.

Method	$q_\alpha(\lambda)$	C'_q	C''_q	v_0	C_v	SVD required
Spectral cut-off	$\frac{1}{\lambda}\mathbf{1}_{[\alpha,\infty)}(\lambda)$	1	1	∞	1	Yes
Tikhonov	$\frac{1}{\lambda+\alpha}$	1	1	1	$v^v(1-v)^{1-v}$	No
m-iterated Tikhonov	$\frac{(\lambda+\alpha)^m-\alpha^m}{\lambda(\lambda+\alpha)^m}$	m	1	m	$(v/m)^v(1-v/m)^{m-v}$	No
Landweber ($\|T\| \le 1$)	$\sum_{j=0}^{\lfloor 1/\alpha\rfloor-1}(1-\lambda)^j$	1	1	∞	$(v/e)^v$	No
Showalter	$\frac{1-\exp\left(-\frac{\lambda}{\alpha}\right)}{\lambda}$	1	1	∞	$(v/e)^v$	No

For further reference, we collect the assumed properties of f and q_α as follows:

Assumption 1. *(i) The true solution f satisfies $f \in \mathcal{W}_\phi$ as in (6).*

(ii) The function ϕ is a qualification of the filter q_α as in (7).

(iii) The function $\psi(x) := x\phi^{-1}(\sqrt{x})$, with ϕ^{-1} being the inverse function of ϕ, is convex.

Remark 2. *We stress that, being a standard assumption for convergence analysis, Assumption 1 (iii) actually imposes no restriction, since one can always work on a slightly larger source set $\mathcal{W}_{\widetilde{\phi}}$ with another index function $\widetilde{\phi}$ for which Assumption 1 (iii) is satisfied.*

2.2 Empirical prediction risk minimization

As discussed in the Introduction, the optimal regularization parameter α will in general depend on Y, σ and f, but the latter is unknown and hence this α cannot be realized in practice. Recall that we always assume $\sigma > 0$ to be known. By meanings of the prediction risk $r(\alpha, f)$, the optimal $\alpha \in \mathcal{A}$ is given by

$$\alpha_o = \operatorname*{argmin}_{\alpha\in\mathcal{A}} r(\alpha, f). \tag{10}$$

As a common remedy, we will try to estimate α_o from the observations Y in (1) by minimizing an unbiased estimator of $r(\alpha, f)$, which can be derived as follows. Let us introduce the shorthand notation

$$s_\alpha(\lambda) := \lambda q_\alpha(\lambda).$$

Then we have

$$\begin{aligned} r(\alpha, f) &= \mathbb{E}\left[\left\|T\left(q_\alpha(T^*T)T^*(Tf+\sigma\xi) - f\right)\right\|_{\mathcal{Y}}^2\right] \\ &= \left\|(I - Tq_\alpha(T^*T)T^*)Tf\right\|_{\mathcal{Y}}^2 + \sigma^2\mathbb{E}\left[\left\|Tq_\alpha(T^*T)T^*\xi\right\|_{\mathcal{Y}}^2\right] \\ &= \sum_{k=1}^{\infty}\lambda_k(1-s_\alpha(\lambda_k))^2 f_k^2 + \sigma^2\sum_{k=1}^{\infty}s_\alpha(\lambda_k)^2. \end{aligned} \tag{11}$$

and furthermore

$$\begin{aligned} &\mathbb{E}\left[\left\|T\hat{f}_\alpha\right\|_{\mathcal{Y}}^2 - 2\langle T\hat{f}_\alpha, Y\rangle_{\mathcal{Y}}\right] \\ =&\,\mathbb{E}\left[\left\|Tq_\alpha(T^*T)T^*(Tf+\sigma\xi)\right\|_{\mathcal{Y}}^2 - 2\langle Tq_\alpha(T^*T)T^*(Tf+\sigma\xi), Tf+\sigma\xi\rangle_{\mathcal{Y}}\right] \\ =&\sum_{k=1}^{\infty}\lambda_k(1-s_\alpha(\lambda_k))^2 f_k^2 - \sum_{k=1}^{\infty}\lambda_k f_k^2 + \sigma^2\sum_{k=1}^{\infty}s_\alpha(\lambda_k)^2 - 2\sigma^2\sum_{k=1}^{\infty}s_\alpha(\lambda_k). \end{aligned}$$

Consequently for $\hat{r}(\alpha, Y)$ as in (3) we have

$$\mathbb{E}\left[\hat{r}(\alpha, Y)\right] = r(\alpha, f) - \sum_{k=1}^{\infty}\lambda_k f_k^2,$$

i.e. up to a constant independent of α, $\hat{r}(\alpha, Y)$ is an unbiased estimator of $r(\alpha, f)$. Hence we define

$$\alpha_{\text{pred}} = \operatorname*{argmin}_{\alpha\in\mathcal{A}}\hat{r}(\alpha, Y). \tag{12}$$

Remark 3. *For the clarity of our notation, we stress that $\alpha \mapsto r(\alpha, f)$ is a deterministic function, whereas $\alpha \mapsto \hat{r}(\alpha, Y)$ is a random function and consequently, α_{pred} as in (11) is a random variable. To obtain bounds for the estimator $\hat{f}_{\alpha_{\text{pred}}}$ we will also need bounds for $r(\alpha_{\text{pred}}, f)$, i.e. a deterministic function evaluated at a random variable. In particular we stress that $\mathbb{E}\left[r(\alpha_{\text{pred}}, f)\right] \neq \mathbb{E}\left[\left\|T\hat{f}_{\alpha_{\text{pred}}} - Tf\right\|_{\mathcal{Y}}^2\right]$.*

Note that α_{pred} is computable in practice as it only relies on the data and the forward operator. As discussed in the Introduction, there are many results available for the performance of $\hat{f}_{\alpha_{\text{pred}}}$ measured by the prediction risk, among which the most precise one is due to Kneip (1994) given below.

Theorem 2.1 (Deviation bound of prediction risk (Kneip, 1994))**.** *Assume the model (1). Let $\hat{f}_\alpha = q_\alpha(T^*T)T^*Y$ with an ordered filter q_α, α_o as in (10), and α_{pred} as in (12). Then there exist positive constants C_1, C_2 such that for all $x \geq 0$ and for all $f \in \mathcal{X}$,*

$$\mathbb{P}\left\{\left\|T\hat{f}_{\alpha_{\text{pred}}} - Tf\right\|_{\mathcal{Y}}^2 - \left\|T\hat{f}_{\alpha_o} - Tf\right\|_{\mathcal{Y}}^2 \geq x\right\} \leq C_1\exp\left(-\frac{C_2}{\sigma}\min\left\{\sqrt{x},\ \frac{x}{\sqrt{r(\alpha_o, f)}}\right\}\right),$$

which remains true when replacing $\left\|T\hat{f}_{\alpha_{\text{pred}}} - Tf\right\|_{\mathcal{Y}}^2 - \left\|T\hat{f}_{\alpha_o} - Tf\right\|_{\mathcal{Y}}^2$ by $r(\alpha_{\text{pred}}, f) - r(\alpha_o, f)$.

Proof. Recall that $\dim \mathcal{X} = \infty$, which, together with the injectivity of T, implies $\dim \mathcal{Y} = \infty$. In fact, the assertion for finite dimensional $\mathcal{Y}$ follows directly from (Kneip, 1994, Proposition 1(i)) and (Kneip, 1994, Theorem 1) by chasing the dependency of the constants on the noise level σ. Concerning the key technical tools in Kneip's proof, we note, that the Lemma 2 there actually holds for infinite sequences $a \in \ell^2$ and bounded linear trace operators $A : \ell^2 \to \ell^2$, and that the Lemma 3 there can be extended to infinite dimensional ordered linear smoothers, as long as they are Hilbert-Schmidt. Thus, the proof by Kneip (1994) carries over to the case that $\dim \mathcal{Y} = \infty$. □

The above theorem, in particular, implies (cf. (Chernousova and Golubev, 2014, Theorem 1))

$$\mathbb{E}\left[\left\|T\hat{f}_{\alpha_{\text{pred}}} - Tf\right\|_{\mathcal{Y}}^2\right] \leq r(\alpha_o, f) + C\sigma\sqrt{r(\alpha_o, f)} \qquad \text{for every } f \in \mathcal{X},$$

which guarantees the order optimality of $\hat{f}_{\alpha_{\text{pred}}}$ in terms of the prediction risk.

3 MISE estimates

The section is devoted to the convergence analysis of $\hat{f}_\alpha = q_\alpha(T^*T)T^*Y$ with $\alpha = \alpha_{\text{pred}}$ as in (12). In what follows, we will prove that $\hat{f}_{\alpha_{\text{pred}}}$ also possesses an optimality property in terms of the direct risk, i.e. in terms of the mean integrated square error (MISE).

3.1 Assumptions and merit discussions

We start with some technical assumptions. As we have already seen in the previous section, many calculations involve summations over $\{\lambda_k\}_{k=1}^\infty$, which can be formulated as Lebesgue-Stieltjes integrals with respect to

$$\Sigma(x) := \#\{k : \lambda_k \geq x\}.$$

Following Bissantz et al. (2007), we assume that Σ can be approximated by a smooth function S to avoid the difficulty caused by the non-smoothness of Σ:

Assumption 2. *(i) There exists a surrogate function $S \in \mathcal{C}^2((0,\infty))$ of Σ satisfying*

$$\begin{aligned}
&\lim_{\alpha \searrow 0} S(\alpha)/\Sigma(\alpha) = 1,\\
&S'(\alpha) < 0 \qquad \text{for all } \alpha > 0,\\
&\lim_{\alpha \nearrow \infty} S(\alpha) = \lim_{\alpha \nearrow \infty} S'(\alpha) = 0,\\
&\lim_{\alpha \searrow 0} \alpha S(\alpha) = 0.
\end{aligned}$$

(ii) There exist constants $\alpha_1 \in (0, \lambda_1]$ and $C_S \in (0,2)$ such that

$$\begin{aligned}
&-\alpha S'(\alpha) \text{ is integrable on } (0, \alpha_1],\\
&\frac{S''(\alpha)}{-S'(\alpha)} \leq \frac{C_S}{\alpha} \qquad \text{for all } \alpha \in (0, \alpha_1].
\end{aligned}$$

(iii) There exists a constant $C_q > 0$ such that

$$\int_1^\infty \Psi'(C_q x) \exp\left(-C_2\sqrt{\frac{x}{2}}\right) \mathrm{d}x < \infty \qquad \text{with } \Psi(x) := \frac{x}{\left(S^{-1}(x)\right)^2}.$$

Here C_2 is as in Theorem 2.1.

Remark 4. *Note that Assumption 2 (i) and (ii) are essentially (Bissantz et al., 2007, Assumption 2) with a proper extension of S to a larger domain $(0, \infty)$. They are rather weak requirements, and cover a wide range of situations, see e.g. (Bissantz et al., 2007, Section 5). The additional assumption (iii) is needed to control certain general moments of $r(\alpha_{\text{pred}}, f)$ with the help of Theorem 2.1.*

The second assumption concerns the choice of regularization parameter α for (ordered) filters q_α.

Assumption 3. *(i) For $\alpha \in \mathcal{A}$, the function $\lambda \mapsto s_\alpha(\lambda)$ is non-decreasing.*

(ii) There exists $c_q > C_q^{-1/2}$ with C_q as in Assumption 2 (iii) such that

$$s_\alpha(\alpha) = \alpha q_\alpha(\alpha) \geq c_q \qquad \text{as } \alpha \searrow 0. \tag{13}$$

Remark 5. *It is easy to see that the above condition with some $c_q > 0$ is satisfied by all the regularization methods in Table 1, which indicates that Assumption 3 is fairly general. More precisely, it can be understood as a parametrization condition. For instance, the Tikhonov method with re-parametrization $\alpha \mapsto \sqrt{\alpha}$, i.e. $q_\alpha(\lambda) = 1/(\sqrt{\alpha} + \lambda)$, still defines an ordered filter, but does not satisfy* (13) *anymore. The condition $c_q > C_q^{-1/2}$ requires some compatibility of the ill-posedness with the parametrization of the filter.*

Under these assumptions, upper bounds for $\hat{f}_\alpha$ with an a-priori choice of α have been proven by Bissantz et al. (2007):

Theorem 3.1 (A-priori parameter choice (Bissantz et al., 2007))**.** *Consider the model* (1)*, and let $\hat{f}_\alpha := q_\alpha(T^*T)T^*Y$ with a filter q_α, and Assumptions 1 (i) and 2 (i)-(ii) hold. Let also α_* satisfy*

$$\alpha_* \phi(\alpha_*)^2 = \sigma^2 S(\alpha_*). \tag{14}$$

(i) If Assumption 1 (ii) holds, then for some constant C_1 we have

$$\sup_{f \in \mathcal{W}_\phi} \mathbb{E}\left[\left\|\hat{f}_{\alpha_*} - f\right\|_{\mathcal{X}}^2\right] \leq C_1 \phi(\alpha_*)^2 = C_1 \sigma^2 \frac{S(\alpha_*)}{\alpha_*} \qquad \text{as } \sigma \searrow 0.$$

(ii) If $\tilde{\phi}(t) := \sqrt{t}\phi(t)$ is a qualification of the filter q_α, then for some constant C_2 we have

$$\sup_{f \in \mathcal{W}_\phi} \mathbb{E}\left[\left\|T\hat{f}_{\alpha_*} - Tf\right\|_{\mathcal{Y}}^2\right] \leq C_2 \alpha_* \phi(\alpha_*)^2 = C_2 \sigma^2 S(\alpha_*) \qquad \text{as } \sigma \searrow 0.$$

3.2 Oracle inequality

We are now in position to derive a crucial oracle inequality in the general form of (4) for the empirical parameter choice $\alpha = \alpha_{\text{pred}}$.

Theorem 3.2 (Oracle inequality)**.** *Assume the model* (1)*. Let $\hat{f}_\alpha := q_\alpha(T^*T)T^*Y$ with an ordered filter q_α, and Assumptions 1, 2 and 3 hold. Let also α_o be given by* (10)*, and α_{pred} by* (12)*. Then there are positive constants C_i, $i = 1, \ldots, 6$, such that for all $f \in \mathcal{W}_\phi$ it holds*

$$\mathbb{E}\left[\left\|\hat{f}_{\alpha_{\text{pred}}} - f\right\|_{\mathcal{X}}^2\right] \leq C_1 \psi^{-1}\Big(C_2 r(\alpha_o, f) + C_3 \sigma^2\Big) + C_4 \frac{r(\alpha_o, f) + \sigma\sqrt{r(\alpha_o, f)}}{S^{-1}\left(C_5 \frac{r(\alpha_o, f)}{\sigma^2}\right)} + C_6 \sigma^2. \tag{15}$$

The proof is based on a general strategy based on technical lemmas, and is given in Section 4. To ease the understanding of this paper, we will now start with conclusions from Theorem 3.2.

3.3 Convergence rates and examples

The derived oracle inequality in Theorem 3.2 readily provides error estimates for the estimator $\hat{f}_{\alpha_{\text{pred}}}$ given proper upper bounds of the oracle prediction risk $r(\alpha_o, f)$ as in Theorem 3.1.

Theorem 3.3 (Convergence rates)**.** *Assume the same setting as in Theorem 3.2, and additionally that $\tilde{\phi}(t) = \sqrt{t}\phi(t)$ is a qualification of the filter q_α. Let α_* be given by (14). Then as $\sigma \searrow 0$ we have*

$$\sup_{f \in \mathcal{W}_\phi} \mathbb{E}\left[\left\|\hat{f}_{\alpha_{\text{pred}}} - f\right\|_{\mathcal{X}}^2\right] \leq C_1 \phi(\alpha_*)^2 + C_2 \frac{\alpha_* \phi(\alpha_*)^2}{S^{-1}\left(C_3 \frac{\alpha_* \phi(\alpha_*)^2}{\sigma^2}\right)} = C_1 \sigma^2 \frac{S(\alpha_*)}{\alpha_*} + C_2 \frac{\sigma^2 S(\alpha_*)}{S^{-1}\left(C_3 S(\alpha_*)\right)}$$

for some positive constants C_i, $i = 1, 2, 3$.

Remark 6. *If, in addition, there is a constant $C_4 > 0$ such that*

$$S(C_4 x) \geq C_3 S(x) \qquad \text{for all } x > 0,$$

then it follows from Theorem 3.3 that

$$\sup_{f \in \mathcal{W}_\phi} \mathbb{E}\left[\left\|\hat{f}_{\alpha_{\text{pred}}} - f\right\|_{\mathcal{X}}^2\right] \leq C \phi(\alpha_*)^2 = C \sigma^2 \frac{S(\alpha_*)}{\alpha_*}.$$

This coincides with the convergence rate we obtain under an a-priori parameter choice in Theorem 3.1, which turns out to be order optimal in most cases, see e.g. Section 3.3.1 and Bissantz et al. (2007).

We note that there are two additional important assumptions for a-posteriori parameter choice α_{pred} compared to the a-priori choice (cf. Theorem 3.1 (i) and Theorem 3.3). The one is Assumption 2 (iii), which concerns the control of general moments of $r(\alpha_{\text{pred}}, f)$, due to the randomness of α_{pred}. The other is that not only ϕ but also $\tilde{\phi}(t) = \sqrt{t}\phi(t)$ is a qualification of the filter q_α. The latter seems to be typical for parameter choice rule relying on residuals in image space, as e.g. for the discrepancy principle in case of deterministic inverse problems, see (Engl et al., 1996, Section 4.3) and (Mathé and Pereverzev, 2006), or for generalized cross-validation (GCV), see (Lukas, 1993; Vogel, 2002). Still we stress that the qualification assumption plays no role in the proof of the oracle inequality (Theorem 3.2), and it only kicks in for the derivation of convergence rates for $\hat{f}_{\alpha_{\text{pred}}}$ through $r(\alpha_o, f)$ in the convergence analysis under the a-priori parameter choice.

3.3.1 Mildly ill-posed problems

We now consider a particular mildly ill-posed problem. More precisely, we assume

Assumption 4. *(i) polynomial decay of eigenvalues of T^*T*

$$\lambda_k = C_a k^{-a} \qquad k = 1, 2, \ldots \qquad \text{with some } a > 1 \text{ and } C_a > 0;$$

(ii) smoothness of the truth

$$f \in \mathcal{W}_{b/(2a)} := \left\{ f \in \mathcal{X} : \sum_{k=1}^{\infty} w_k f_k^2 \leq 1 \right\} \qquad \text{with } w_k = C_b k^b \quad k = 1, 2, \ldots,$$

for some positive constants b and C_b.

In the above assumption, the requirement of $a > 1$ is to ensure that the forward operator T is Hilbert-Schmidt, and the notation $\mathcal{W}_{b/(2a)}$ is consistent with (8) as the above smoothness condition is equivalent to the Hölder type source condition of order $b/(2a)$.

In this simple setting the convergence rates of ordered filter based methods with empirical parameter choice α_{pred} can be explicitly computed.

Corollary 3.4 (Order optimality)**.** *Assume the model* (1)*. Let* $\hat{f}_\alpha := q_\alpha(T^*T)T^*Y$ *with an ordered filter* q_α*, and Assumptions 3 and 4 with any* $c_q > 0$ *hold. If the qualification index* v_0 *in* (9) *of the ordered filter satisfies* $v_0 \geq b/(2a) + 1/2$*, then there are positive constants* C_1 *and* C_2 *such that*

$$\sup_{f\in\mathcal{W}_{b/(2a)}} \mathbb{E}\left[\left\|\hat{f}_{\alpha_{\text{pred}}} - f\right\|_{\mathcal{X}}^2\right] \leq C_1 \inf_{\hat{f}} \sup_{f\in\mathcal{W}_{b/(2a)}} \mathbb{E}\left[\left\|\hat{f} - f\right\|_{\mathcal{X}}^2\right] \leq C_2\sigma^{\frac{2b}{a+b+1}} \qquad \text{as } \sigma \searrow 0.$$

The infimum above is taken over all possible estimators $\hat{f}$ *(including both linear and nonlinear ones).*

Proof. Define $S(\alpha) := \left(\alpha/C_a\right)^{-1/a}$, which satisfies Assumption 2 (i) and (ii). Note that $\phi(x) = x^{b/(2a)}$ and then $\psi(x) := x\phi^{-1}(x) = x^{(a+b)/b}$, so Assumption 1 is satisfied. Since $\Psi(x) = x/\left(S^{-1}(x)\right)^2 = C_a^{-2}x^{1+2a}$, Assumption 2 (iii) clearly holds. By definition (14), it follows that $\alpha_* = C_a^{1/(1+a+b)}\sigma^{2a/(1+a+b)}$. Thus, by Theorem 3.3 and Remark 6, we have

$$\sup_{f\in\mathcal{W}_{b/(2a)}} \mathbb{E}\left[\left\|\hat{f}_{\alpha_{\text{pred}}} - f\right\|_{\mathcal{X}}^2\right] \leq C_1\phi(\alpha_*)^2 = C_1 C_a^{\frac{b}{a(1+a+b)}}(\sigma^2)^{\frac{b}{1+a+b}} \qquad \text{as } \sigma \searrow 0,$$

for some positive constant C_1.

Moreover, it is well-known that (see e.g. Pinsker, 1980; Golubev and Khas'minskiĭ, 1999)

$$\inf_{\hat{f}} \sup_{f\in\mathcal{W}_{b/(2a)}} \mathbb{E}\left[\left\|\hat{f} - f\right\|_{\mathcal{X}}^2\right] \geq C_2(\sigma^2)^{\frac{b}{a+b+1}} \qquad \text{for some } C_2 > 0.$$

□

Remark 7. *The above proposition in particular implies that all the methods in Table 1 are order optimal with* $\alpha = \alpha_{\text{pred}}$ *when* $b/(2a) + 1/2 \leq v_0$ *in the minimax sense. This reproduces the result in Chernousova and Golubev (2014) for the spectral cut-off method as a special case.*

Note that the parameter choice α_{pred} *depends only on the data, and is completely independent of the unknown truth. Thus, the ordered filter based regularization methods with* α_{pred} *automatically adapt to the unknown smoothness of the truth, and achieve the best possible rates up to a constant. In other words, the ordered filter based regularization methods with parameter choice* α_{pred} *are* adaptively minimax optimal *(cf. Pinsker, 1980) over a range of smoothness classes, and the adaptation range is determined by the qualification index of the filter and the smoothing property of the forward operator. Importantly, we point out that the order optimality here is in sharp contrast to the Lepskiĭ-type balancing principle (Lepskiĭ, 1991), where one typically losses a log-factor in the asymptotic convergence rates.*

In addition, we stress again the price we pay for a-posteriori parameter choice α_{pred} *is a stronger qualification assumption* $b/(2a) + 1/2 \leq v_0$*, as the convergence rates of direct risk for a-priori parameter choice only asks for* $b/(2a) \leq v_0$ *by Theorem 3.1 (i), see also Remark 6.*

3.3.2 Exponentially ill-posed problems

Next we consider an exponentially ill-posed setting: The eigenvalues of T^*T satisfy

$$\lambda_k \asymp \exp\left(-\mu k^a\right) \qquad k = 1, 2, \ldots \qquad \text{for some } a, \mu > 0,$$

and the smoothness of the truth is characterized by (6) with

$$\phi(x) := x^b(-\log x)^c \qquad \text{for some } b > 0, c \in \mathbb{R}, \text{ or } b = 0, c > 0.$$

That is, we assume that the truth lies in $\left\{f \in \mathcal{X} : \sum_{k=1}^\infty w_k f_k^2 \leq 1\right\}$ with $w_k \asymp k^{-2ac}\exp(2\mu b k^a)$.

It is easy to see that the assumptions of Theorem 3.2 are satisfied for all the regularization methods listed in Table 1 provided that $0 < a < 1/2$, or $a = 1/2$ and μ is sufficiently small. Hence, if in addition $b + 1/2 < v_0$, or $b = v_0, c \leq 0$ (recall that v_0 is the qualification index),

we can obtain certain error bounds by means of Theorem 3.3. However, it turns out that our bounds are too rough to guarantee order optimality. In fact, in case of $b = 0$, the error bound on the right-hand side of (15) even diverges as the noise level σ tends to 0. In summary, our oracle inequality is applicable for exponentially ill-posed problems as discussed here, but is not strong enough to derive rates of convergence or even show optimality of the investigated parameter choice α_{pred}.

4 A general methodology for proving oracle inequalities of the form (15)

To prove the oracle inequality (15), we will proceed as follows. First, we obtain bounds for general moments of the prediction risk (cf. Corollary 4.1). Secondly, we proceed with a standard estimate for the (deterministic) bias of the estimator $\hat{f}_\alpha$ under the smoothness Assumption 1 (cf. Lemma 4.2). Finally, we prove a comparison Lemma on the variance terms (cf. Lemma 4.3). Putting these three ingredients together, the proof of Theorem 3.2 is then straightforward. Rephrased differently, if these three ingredients are given for some parameter choice rule $\bar{\alpha}$, an oracle inequality similar to (15) can be derived.

4.1 Main ingredients

We start with bounds for general moments of the prediction risk.

Corollary 4.1 (General moments of prediction risk)**.** *Assume the same setting as Theorem 2.1, and let $\Psi : \mathbb{R}_+ \to \mathbb{R}_+$ be an increasing and continuously differentiable function with $\Psi(0) = 0$ such that*

$$\int_1^\infty \Psi'(x) \exp\left(-C_2\sqrt{\frac{x}{2}}\right) dx < \infty \tag{16}$$

with C_2 being the same as in Theorem 2.1. Then

$$\mathbb{E}\left[\Psi\left(\frac{r(\alpha_{\text{pred}}, f)}{\sigma^2}\right)\right] \leq \Psi\left(\frac{2r(\alpha_o, f)}{\sigma^2}\right) + C_\Psi \qquad \text{for all } f \in \mathcal{X},$$

where C_Ψ is a constant depending only on Ψ.

Proof. Let $r_o := r(\alpha_o, f)$. Then

$$\begin{aligned}
\mathbb{E}\left[\Psi\left(\frac{r(\alpha_{\text{pred}}, f)}{\sigma^2}\right)\right] &= \int_0^\infty \Psi'(x)\mathbb{P}\left\{\frac{r(\alpha_{\text{pred}}, f)}{\sigma^2} \geq x\right\} \mathrm{d}x && \text{[by Fubini's theorem]} \\
&\leq \int_0^{2r_o/\sigma^2} \Psi'(x)\,\mathrm{d}x + C_1 \int_{2r_o/\sigma^2}^\infty \Psi'(x) \exp\left(-C_2\sqrt{x - \frac{r_o}{\sigma^2}}\right) \mathrm{d}x && \text{[by Theorem 2.1]} \\
&\leq \Psi\left(\frac{2r_o}{\sigma^2}\right) + C_1 \int_{2r_o/\sigma^2}^\infty \Psi'(x) \exp\left(-C_2\sqrt{\frac{x}{2}}\right) \mathrm{d}x \\
&\leq \Psi\left(\frac{2r(\alpha_o, f)}{\sigma^2}\right) + C_\Psi.
\end{aligned}$$

This concludes the proof. □

Remark 8. *A simple example is $\Psi(x) := x^\theta$ for some $\theta > 0$. By Corollary 4.1, it leads to*

$$\mathbb{E}\left[r(\alpha_{\text{pred}}, f)^\theta\right] \leq 2^\theta r(\alpha_o, f)^\theta + C_\theta \sigma^{2\theta},$$

where the second term on the right hand side is typically negligible compared to the first as σ tends to 0.

Based on the smoothness Assumption 1, we can prove the following estimate for the bias of $\hat{f}_\alpha$:

Lemma 4.2 (Source condition)**.** *Under Assumption 1 (i) and (iii), it holds that*

$$\sum_{k=1}^{\infty}(1-s_\alpha(\lambda_k))^2 f_k^2 \le C^2\psi^{-1}\left(\frac{1}{C^2}\sum_{k=1}^{\infty}\lambda_k(1-s_\alpha(\lambda_k))^2 f_k^2\right),$$

with C being the same constant in (6)*.*

Proof. Based on spectral analysis, condition (6) can be equivalently written as

$$f_k = \phi(\lambda_k)w_k, \quad k=1,2,\ldots, \qquad \text{with } \sum_{k=1}^{\infty} w_k^2 \le C^2.$$

Thus, $f_k^2 \le C^2\phi(\lambda_k)^2$ and then $\phi^{-1}(\sqrt{C^{-2}f_k^2}) \le \lambda_k$. By Jensen's inequality, we have

$$\begin{aligned}\psi\Big(\sum_{k=1}^{\infty}\frac{1}{C^2}f_k^2\left(1-s_\alpha(\lambda_k)\right)^2\Big) &\le \sum_{k=1}^{\infty}\psi\Big(\frac{1}{C^2}f_k^2\Big)\left(1-s_\alpha(\lambda_k)\right)^2\\ &\le \frac{1}{C^2}\sum_{k=1}^{\infty}\lambda_k f_k^2\left(1-s_\alpha(\lambda_k)\right)^2.\end{aligned}$$

The assertion follows by applying ψ^{-1} to the above inequality. □

Finally, we need certain comparison relations, which are required for bounding the variance part of the risk:

Lemma 4.3 (Comparison)**.** *Let q_α be a filter, and Assumptions 2 and 3 hold. Then there are positive constants C_1 and C_2 such that for every $\alpha \in \mathcal{A}$*

$$\sum_{k=1}^{\infty}\lambda_k q_\alpha(\lambda_k)^2 \le C_1\Psi_1\left(C_q\sum_{k=1}^{\infty}s_\alpha(\lambda_k)^2\right) \quad \text{and} \quad \sum_{k=1}^{\infty}\lambda_k^2 q_\alpha(\lambda_k)^4 \le C_2\Psi_2\left(C_q\sum_{k=1}^{\infty}s_\alpha(\lambda_k)^2\right),$$

with C_q being the same as in Assumption 2 (iii), and the functions Ψ_1 and Ψ_2 are defined as

$$\Psi_1(x) = \frac{x}{S^{-1}(x)}, \qquad \Psi_2(x) = \frac{x}{\left(S^{-1}(x)\right)^2}, \qquad x>0. \tag{17}$$

Proof. To ease notation, let C be a generic positive constant, whose value may vary from place to place. We consider two separate cases: α is sufficiently small, and $\alpha \not\searrow 0$.

Case I: α is sufficiently small. Let c_q be the same as in Assumption 3. Then as $\alpha \searrow 0$

$$\begin{aligned}\sum_{k=1}^{\infty}s_\alpha(\lambda_k)^2 &= -\int_0^{\infty}t^2 q_\alpha(t)^2\,\mathrm{d}\Sigma(t) \ge -\int_\alpha^{\infty}t^2 q_\alpha(t)^2\,\mathrm{d}\Sigma(t)\\ &\ge -\int_\alpha^{\infty}\alpha^2 q_\alpha(\alpha)^2\,\mathrm{d}\Sigma(t) && \text{[by Assumption 3 (i)]}\\ &\ge -c_q^2\int_\alpha^{\infty}\mathrm{d}\Sigma(t) = c_q^2\Sigma(\alpha) && \text{[by Assumption 3 (ii)]}\\ &= c_q^2 S(\alpha). && \text{[by Assumption 2 (i)]}\end{aligned} \tag{18}$$

It follows from Assumption 2 (ii) and (Bissantz et al., 2007, Lemma 12) that as $\alpha \searrow 0$

$$\int_0^{\alpha}S(t)\,\mathrm{d}t = \alpha S(\alpha) - \int_0^{\alpha}t\,\mathrm{d}S(t) \le C\alpha S(\alpha) - C\alpha^2\int_\alpha^{\alpha_1}\frac{1}{t}\,\mathrm{d}S(t) \le C\alpha S(\alpha). \tag{19}$$

Then as $\alpha \searrow 0$ we have

$$\begin{aligned}
\sum_{k=1}^{\infty} \lambda_k q_\alpha(\lambda_k)^2 &= -\int_0^\infty t q_\alpha(t)^2 \,\mathrm{d}\Sigma(t) \\
&\le C\left(-\int_0^\alpha \frac{t}{\alpha^2}\,\mathrm{d}\Sigma(t) - \int_\alpha^\infty \frac{1}{t}\,\mathrm{d}\Sigma(t)\right) && \text{[by Definition 1 (i)]} \\
&\le C\left(-\frac{t}{\alpha^2}\Sigma(t)\Big|_0^\alpha + \frac{1}{\alpha^2}\int_0^\alpha \Sigma(t)\,\mathrm{d}t + \frac{1}{\alpha}\Sigma(\alpha)\right) = C\frac{1}{\alpha^2}\int_0^\alpha \Sigma(t)\,\mathrm{d}t \\
&= C\frac{1}{\alpha^2}\int_0^\alpha S(t)\,\mathrm{d}t && \text{[by Assumption 2 (i)]} \\
&\le C\frac{S(\alpha)}{\alpha}, && \text{[by (19)]}
\end{aligned}$$

and similarly

$$\begin{aligned}
\sum_{k=1}^{\infty} \lambda_k^2 q_\alpha(\lambda_k)^4 &= -\int_0^\infty t^2 q_\alpha(t)^4 \,\mathrm{d}\Sigma(t) \\
&\le C\left(-\int_0^\alpha \frac{t^2}{\alpha^4}\,\mathrm{d}\Sigma(t) - \int_\alpha^\infty \frac{1}{t^2}\,\mathrm{d}\Sigma(t)\right) && \text{[by Definition 1 (i)]} \\
&\le C\left(-\frac{t^2}{\alpha^4}\Sigma(t)\Big|_0^\alpha + \frac{1}{\alpha^4}\int_0^\alpha 2t\Sigma(t)\,\mathrm{d}t + \frac{1}{\alpha^2}\Sigma(\alpha)\right) \le C\frac{1}{\alpha^3}\int_0^\alpha \Sigma(t)\,\mathrm{d}t \\
&= C\frac{1}{\alpha^3}\int_0^\alpha S(t)\,\mathrm{d}t && \text{[by Assumption 2 (i)]} \\
&\le C\frac{S(\alpha)}{\alpha^2}. && \text{[by (19)]}
\end{aligned}$$

These together with (18) prove the assertion if α is sufficiently small, by noticing that $C_q > 1/c_q^2$.

Case II: $\alpha \not\searrow 0$. Note that $\alpha \not\searrow 0$ and the boundedness of $\mathcal{A}$ imply that $1/\alpha$ is bounded below from 0, and above from ∞. Then by Assumption 3 it holds that

$$\sum_{k=1}^{\infty} s_\alpha(\lambda_k)^2 \ge C\frac{1}{\alpha}\sum_{k=1}^{\infty}\lambda_k \ge C.$$

By Definition 1 (i) we have

$$\sum_{k=1}^{\infty} \lambda_k q_\alpha(\lambda_k)^2 \le C\frac{1}{\alpha^2}\sum_{k=1}^{\infty}\lambda_k \le C, \quad \text{and} \quad \sum_{k=1}^{\infty} \lambda_k^2 q_\alpha(\lambda_k)^4 \le C\frac{1}{\alpha^3}\sum_{k=1}^{\infty}\lambda_k \le C.$$

Thus the assertion clearly holds for $\alpha \not\searrow 0$.

Combining the above two cases concludes the proof. □

4.2 Proof of Theorem 3.2

Now we are in position to prove Theorem 3.2. To obtain bounds, we will split the variance part by exploiting a technique from Golubev (2010) dealing with ordered processes:

Lemma 4.4 (Ordered processes (Golubev, 2010))**.** *Assume that a sequence of functions $c_k : \mathcal{A} \to \mathbb{R}$, $k = 1, 2, \ldots$, with $\mathcal{A} \subseteq \mathbb{R}_+$, satisfy that*

$$\kappa(\alpha) := \left(\sum_{k=1}^{\infty} c_k(\alpha)^2\right)^{1/2} < \infty \qquad \textit{for } \alpha \in \mathcal{A},$$

and that κ is continuous and strictly monotone on $\mathcal{A}$. Define $(x)_+ := \max\{x, 0\}$, and

$$\zeta(\alpha) := \sum_{k=1}^{\infty} c_k(\alpha)(\xi_k^2 - 1) \qquad \text{with } \xi_k \overset{iid}{\sim} \mathcal{N}(0,1).$$

Then there exists a constant C such that

$$\mathbb{E}\left[\sup_{\alpha\in\mathcal{A}} \left(\zeta(\alpha) - x\kappa(\alpha)^2\right)_+\right] \leq \frac{C}{x} \qquad \text{for all } x > 0.$$

Proof. Note that $\zeta(\alpha)$ almost surely finite since $\kappa(\alpha) < \infty$. In the terminology of Golubev (2010), $\zeta(\alpha)$ is an ordered process, but with a possibly nonzero mean. The assertion can be proven in exactly the same way as (Golubev, 2010, Lemmas 1 and 2). □

Proof of Theorem 3.2. Consider the bias-variance decomposition

$$\begin{aligned}
\mathbb{E}\left[\left\|\hat{f}_{\alpha_{\text{pred}}} - f\right\|_{\mathcal{X}}^2\right] &\leq 2\mathbb{E}\left[\left\|(q_{\alpha_{\text{pred}}}(T^*T)T^*T - I)f\right\|_{\mathcal{X}}^2\right] + 2\sigma^2\mathbb{E}\left[\left\|q_{\alpha_{\text{pred}}}(T^*T)T^*\xi\right\|_{\mathcal{X}}^2\right] \\
&= 2\mathbb{E}\left[\sum_{k=1}^{\infty}\left(1 - s_{\alpha_{\text{pred}}}(\lambda_k)\right)^2 f_k^2\right] + 2\sigma^2\mathbb{E}\left[\sum_{k=1}^{\infty}\lambda_k q_{\alpha_{\text{pred}}}(\lambda_k)^2\xi_k^2\right]. \qquad (20)
\end{aligned}$$

For the first term (i.e. bias part) in (20), we have for all $f \in \mathcal{W}$

$$\begin{aligned}
\mathbb{E}\left[\sum_{k=1}^{\infty} f_k^2\left(1 - s_{\alpha_{\text{pred}}}(\lambda_k)\right)^2\right] &\leq C_1\mathbb{E}\left[\psi^{-1}\Big(\frac{1}{C_1}\sum_{k=1}^{\infty}\lambda_k(1 - s_{\alpha_{\text{pred}}}(\lambda_k))^2 f_k^2\Big)\right] && \text{[by Lemma 4.2]} \\
&\leq C_1\mathbb{E}\left[\psi^{-1}\Big(\frac{1}{C_1}r(\alpha_{\text{pred}}, f)\Big)\right] && \text{[by (11)]} \\
&\leq C_1\psi^{-1}\Big(\frac{1}{C_1}\mathbb{E}\left[r(\alpha_{\text{pred}}, f)\right]\Big) && \text{[by Jensen's inequality]} \\
&\leq C_1\psi^{-1}\Big(\frac{2}{C_1}r(\alpha_o, f) + C_2\sigma^2\Big). && \text{[by Remark 8]}
\end{aligned}$$

For the second term (i.e. variance part) in (20), we further split it into two terms

$$\sigma^2\mathbb{E}\left[\sum_{k=1}^{\infty}\lambda_k q_{\alpha_{\text{pred}}}(\lambda_k)^2\xi_k^2\right] = \sigma^2\mathbb{E}\left[\sum_{k=1}^{\infty}\lambda_k q_{\alpha_{\text{pred}}}(\lambda_k)^2\right] + \sigma^2\mathbb{E}\left[\sum_{k=1}^{\infty}\lambda_k q_{\alpha_{\text{pred}}}(\lambda_k)^2(\xi_k^2 - 1)\right]. \qquad (21)$$

Note that $\Psi_1'(x) \leq S^{-1}(x)\Psi_2'(x)$ for the functions Ψ_1 and Ψ_2 in (17), so by monotonicity of S and Assumption 2 (iii), condition (16) is satisfied with $\Psi(x) = \Psi_1(C_q x)$ as well. Then, for the first term in (21), it holds that

$$\begin{aligned}
\sigma^2\mathbb{E}\left[\sum_{k=1}^{\infty}\lambda_k q_{\alpha_{\text{pred}}}(\lambda_k)^2\right] &\leq C_3\sigma^2\mathbb{E}\left[\Psi_1\left(C_q\sum_{k=1}^{\infty}s_{\alpha_{\text{pred}}}(\lambda_k)^2\right)\right] && \text{[by Lemma 4.3]} \\
&\leq C_3\sigma^2\mathbb{E}\left[\Psi_1\left(C_q\frac{r(\alpha_{\text{pred}}, f)}{\sigma^2}\right)\right] && \text{[by (11)]} \\
&\leq C_3\sigma^2\left(\Psi_1\left(2C_q\frac{r(\alpha_o, f)}{\sigma^2}\right) + C_4\right) && \text{[by Corollary 4.1]} \\
&= 2C_3C_q\frac{r(\alpha_o, f)}{S^{-1}\left(2C_q\frac{r(\alpha_o,f)}{\sigma^2}\right)} + C_3C_4\sigma^2.
\end{aligned}$$

For the second term in (21), we apply Lemma 4.4 with $c_k(\alpha) = \lambda_k q_\alpha(\lambda_k)^2$, which leads to (with ζ and κ defined in Lemma 4.4)

$$\begin{aligned}
\sigma^2\mathbb{E}\left[\sum_{k=1}^{\infty}\lambda_k q_{\alpha_{\text{pred}}}(\lambda_k)^2(\xi_k^2-1)\right] &\leq \sigma^2\mathbb{E}\left[\left(\zeta(\alpha_{\text{pred}}) - x\kappa(\alpha_{\text{pred}})^2\right)_+\right] + \sigma^2\mathbb{E}\left[x\kappa(\alpha_{\text{pred}})^2\right] \\
&\leq \sigma^2\mathbb{E}\left[\sup_{\alpha\in\mathcal{A}}\left(\zeta(\alpha) - x\kappa(\alpha)^2\right)_+\right] + \sigma^2\mathbb{E}\left[x\kappa(\alpha_{\text{pred}})^2\right] \\
&\leq \sigma^2\frac{C_5}{x} + x\sigma^2\mathbb{E}\left[\sum_{k=1}^{\infty}\lambda_k^2 q_{\alpha_{\text{pred}}}(\lambda_k)^4\right] \qquad \text{for all } x>0.
\end{aligned}$$

We minimize the right hand side of the above equation over $x > 0$, and then obtain

$$\begin{aligned}
\sigma^2\mathbb{E}\left[\sum_{k=1}^{\infty}\lambda_k q_{\alpha_{\text{pred}}}(\lambda_k)^2(\xi_k^2-1)\right] &\leq 2\sigma^2\left(C_5\mathbb{E}\left[\sum_{k=1}^{\infty}\lambda_k^2 q_{\alpha_{\text{pred}}}(\lambda_k)^4\right]\right)^{1/2} \\
&\leq C_6\sigma^2\left(\mathbb{E}\left[\Psi_2\left(C_q\sum_{k=1}^{\infty}s_{\alpha_{\text{pred}}}(\lambda_k)^2\right)\right]\right)^{1/2} && \text{[by Lemma 4.3]} \\
&\leq C_6\sigma^2\left(\mathbb{E}\left[\Psi_2\left(C_q\frac{r(\alpha_{\text{pred}},f)}{\sigma^2}\right)\right]\right)^{1/2} && \text{[by (11)]} \\
&\leq C_6\sigma^2\left(\Psi_2\left(2C_q\frac{r(\alpha_o,f)}{\sigma^2}\right)+C_7\right)^{1/2} && \text{[by Corollary 4.1]} \\
&\leq C_6(2C_q)^{1/2}\frac{\sigma\sqrt{r(\alpha_o,f)}}{S^{-1}\left(2C_q\frac{r(\alpha_o,f)}{\sigma^2}\right)} + C_6(C_7)^{1/2}\sigma^2.
\end{aligned}$$

Combining all these estimates concludes the proof. □

5 Numerical simulations

In this section, we will investigate the behavior of the following parameter choice methods by means of a simulation study:

(i) The oracle parameter choice $\alpha_{\text{or}} = \operatorname{argmin}_{\alpha\in\mathcal{A}}\mathbb{E}\left[\left\|\hat{f}_\alpha - f^\dagger\right\|\right]$, which is not available in practice and evaluated here for comparison only,

(ii) the a-posteriori parameter choice rule which is investigated in this paper, given by $\alpha_{\text{pred}} = \operatorname{argmin}_{\alpha\in\mathcal{A}}\hat{r}(\alpha, Y)$ with $\hat{r}(\alpha, Y)$ as in (3),

(iii) and the Lepskiĭ-type balancing principle originally introduced by Lepskiĭ (1991), and was further developed for usage in statistical inverse problems by Bauer and Hohage (2005); Mathé (2006); Mathé and Pereverzev (2006); Werner and Hohage (2012). Given estimators $\hat{f}_{\alpha_1}$, ..., $\hat{f}_{\alpha_n}$ for a suitably fine grid $\alpha_1 < ... < \alpha_n$, it consists in choosing

$$j^* = \max\left\{1\leq j\leq n \;\middle|\; \left\|\hat{f}_{\alpha_j} - \hat{f}_{\alpha_k}\right\|_{\mathcal{X}} \leq 4\sigma\sqrt{\text{Trace}\left(q_{\alpha_k}(T^*T)\,T^*\right)} \text{ for all } 1\leq k\leq j\right\}. \tag{22}$$

and $\alpha_{\text{LEP}} := \alpha_{j^*}$. Note that the term $\sigma\sqrt{\text{Trace}\left(q_\alpha(T^*T)\,T^*\right)}$ is in fact an estimator for the standard deviation of $\hat{f}_\alpha$. For an explanatory derivation of this choice we refer to Mathé (2006). Unfortunately, the computation of α_{LEP} is computationally expensive (see e.g. our simulations below).

In all above methods, the set of considered parameters is

$$\mathcal{A} = \left\{ \sigma^2 \cdot r^k \;\middle|\; k = 0, 1, ..., \left\lfloor (\log(r))^{-1} \log\left(\sigma^{-2} \|T^*T\|\right) \right\rfloor \right\}$$

for some $r > 1$, i.e. it discretizes the range of possible α's $\left[\sigma^2, \|T^*T\|\right]$ in a logarithmically equispaced way. In our simulations we use $r = 1.2$. We also tried different values of r which did not influence the results significantly. Note that it can readily be seen from the error decomposition (11) that the discrete parameter set $\mathcal{A}$ is – under appropriate conditions on the filter which are satisfied by all filters in Table 1 – able to resemble the optimal behavior of a continuous parameter set $\sigma^2 \leq \alpha < \infty$ up to a constant depending on r.

Let us briefly comment on the implementation of the parameter choice rule α_{pred}. Even though the minimization is not performed over the continuum $\alpha \in \mathcal{A}$ but over a discrete set here, the computation of α_{pred} can be numerically challenging. In Figure 1 we depict the function $\alpha \mapsto \hat{r}(\alpha, Y)$ in an example using Tikhonov regularization, which shows that the function is quite flat around its absolute minimum. We observed this behavior in many situations, especially if σ is not too small. Finally we mention that the evaluation of the trace operator in $\hat{r}(\alpha, Y)$ can be expensive, but this can be overcome by different techniques (cf. (Engl et al., 1996, Sec. 9.4) or (Vogel, 2002, Sec. 7.1)).

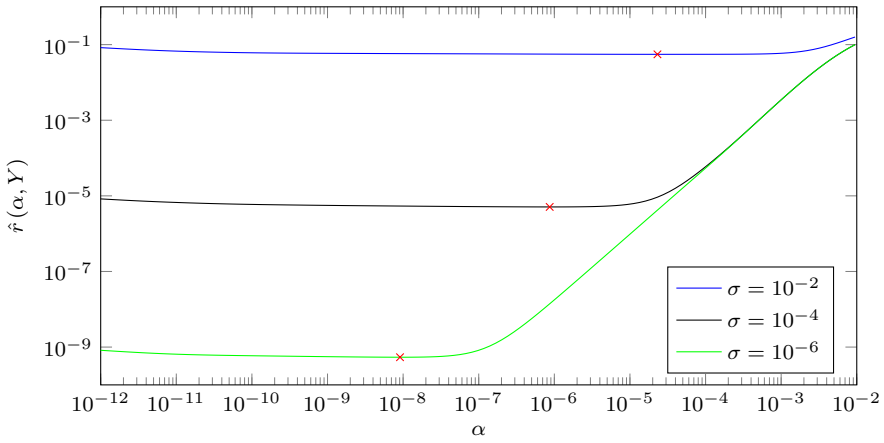

Figure 1: The function $\alpha \mapsto \hat{r}(\alpha, Y)$ for different noise levels. The red crosses mark the actual minima. The operator T is as in Section 5.1, and f is as in Example 1. The chosen regularization method is Tikhonov regularization.

5.1 Convergence rates

At first we investigate the empirical rate of convergence in a mildly ill-posed situation. Therefore we consider the linear integral operator $T : \mathbf{L}^2([0,1]) \to \mathbf{L}^2([0,1])$ defined by

$$(Tf)(x) = \int_0^1 k(x,y) f(y) \, \mathrm{d}y, \qquad x \in [0,1]$$

with kernel $k(x,y) = \min\{x \cdot (1-y), y \cdot (1-x)\}, x, y \in [0,1]$, i.e. $(Tf)'' = -f$ for all $f \in \mathbf{L}^2([0,1])$. Obviously, the eigenvalues λ_k of T^*T satisfy $\lambda_k \sim k^{-4}$.

We discretize T by choosing equidistant points $x_1 = \frac{1}{2n}, x_2 = \frac{3}{2n}, \ldots, x_n = \frac{2n-1}{2n}$ and using the composite midpoint rule

$$(Tf)(x) = \int_0^1 k(x,y) f(y) \, \mathrm{d}y \approx \frac{1}{n} \sum_{i=1}^{n} k(x, x_i) f(x_i)$$

on the grid points $x = x_j$, $1 \leq j \leq n$. To avoid an inverse crime, the exact data Tf is always calculated analytically.

We consider two different scenarios varying in the smoothness of the unknown solution f:

Example 1. *As the first example, we consider the continuous function*

$$f(x) = \begin{cases} x & \text{if } 0 \leq x \leq \frac{1}{2}, \\ 1-x & \text{if } \frac{1}{2} \leq x \leq 1. \end{cases}$$

It can readily be seen by straightforward computations that the Fourier coefficients f_k of f are given by

$$f_k = \frac{(-1)^k - 1}{4\pi^3 k^2}.$$

Consequently $f \in \mathcal{W}_{3/8-\varepsilon}$ and the optimal rate of convergence is $\mathcal{O}\left(\sigma^{\frac{3}{4}-\varepsilon}\right)$ for any $\varepsilon > 0$, see Proposition 3.4.

Example 2. *In the second example we choose*

$$f(x) = \begin{cases} 1 & \text{if } \frac{1}{4} \leq x \leq \frac{3}{4}, \\ 0 & \text{else.} \end{cases}$$

As this function can be written as the derivative of functions as in the first example, it is clear that $f \in H^{\frac{1}{2}-\varepsilon}([0,1])$ for any $\varepsilon > 0$. This is also evident by the fact that the Fourier coefficients f_k of f are given by

$$f_k = \frac{(-1)^k \sin\left(\frac{\pi k}{2}\right)}{2\pi^2 k}.$$

Consequently, $f \in \mathcal{W}_{1/8-\varepsilon}$ and the optimal rate of convergence is $\mathcal{O}\left(\sigma^{\frac{1}{3}-\varepsilon}\right)$ for any $\varepsilon > 0$, see Proposition 3.4.

In Figures 2-4 we plot several empirical risks against the noise level $\sigma \in \left\{2^{-15}, ..., 2^{-25}\right\}$. The optimal rate of convergence is also indicated. We consider spectral cut-off (cf. Figure 2), Tikhonov regularization (cf. Figure 3), and Showalter regularization (cf. Figure 4). Using Monte Carlo simulations with 10^4 experiments per noise level we compute empirical versions of the oracle risk $R_{\text{or}}(\sigma) := \mathbb{E}\left[\left\|\hat{f}_{\alpha_{\text{or}}} - f^\dagger\right\|_2^2\right]$, the prediction risk $R_{\text{pred}}(\sigma) := \mathbb{E}\left[\left\|\hat{f}_{\alpha_{\text{pred}}} - f^\dagger\right\|_2^2\right]$, and the Lepskiĭ risk $R_{\text{LEP}}(\sigma) := \mathbb{E}\left[\left\|\hat{f}_{\alpha_{\text{LEP}}} - f^\dagger\right\|_2^2\right]$.

In all plots we find a good agreement of our theoretical predictions and the empirical results. Compared with the Lepskiĭ-type balancing principle, it seems that α_{pred} performs order-optimal with a slightly smaller constant. The loss of a log-factor by using α_{LEP} cannot be visible in such a small simulation study.

5.2 Efficiency simulations

Besides the convergence rates simulations above we also want to numerically infer on the constant in the oracle inequality which will be done by efficiency simulations. Therefore (inspired by Bauer and Lukas (2011); Chernousova and Golubev (2014)) we consider the following setup. The forward operator is a 300×300 diagonal matrix with singular values $\lambda(k) = k^{-a}$ with a fixed parameter $a > 0$. Then we repeat the following experiment 10^4 times : Given a parameter ν we generate a random ground truth $f \in \mathbb{R}^{300}$ by $f(k) = \pm k^{-\nu} \cdot \left(1 + \mathcal{N}\left(0, 0.1^2\right)\right)$ where the sign is independent and uniformly distributed for each component. From this ground truth, data is generated according to $Y(k) = \lambda(k) \cdot f(k) + \mathcal{N}\left(0, \sigma^2\right)$ where the noise is again independent in each component. Based on the data compute empirical versions of the oracle risk $R_{\text{or}}(\sigma)$, the

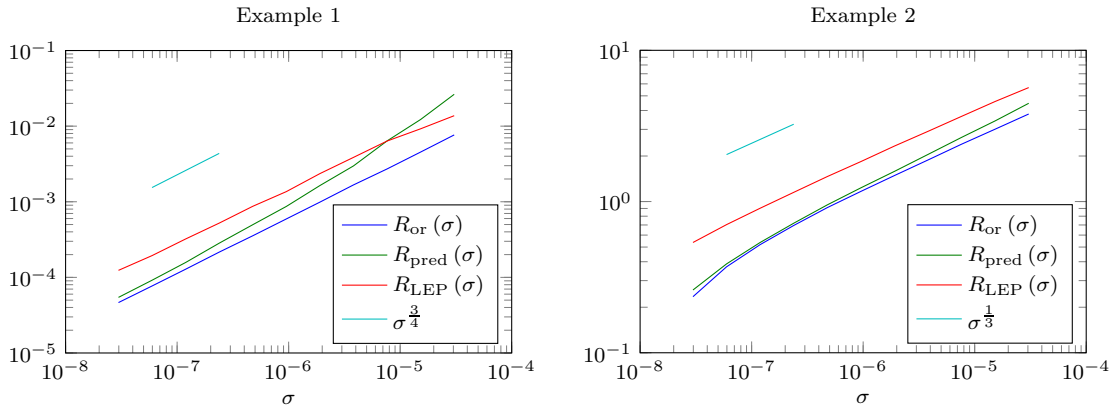

Figure 2: Simulation results for spectral cut-off.

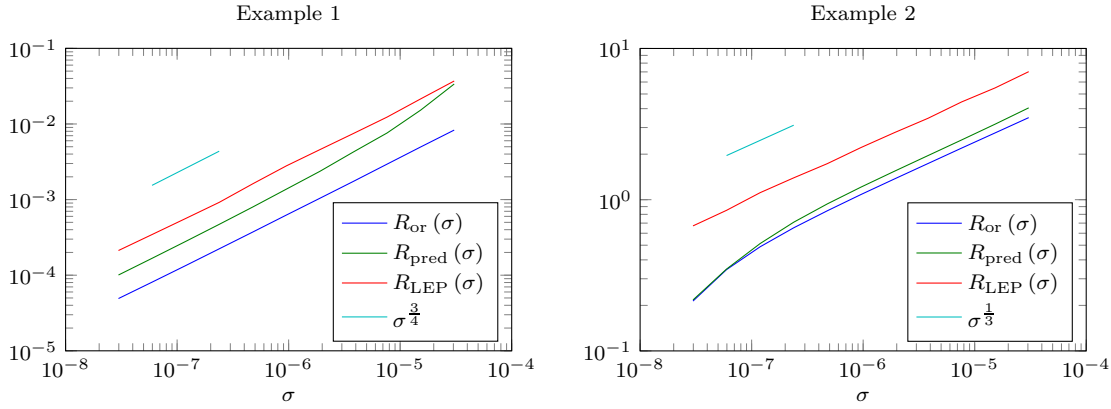

Figure 3: Simulation results for Tikhonov regularization.

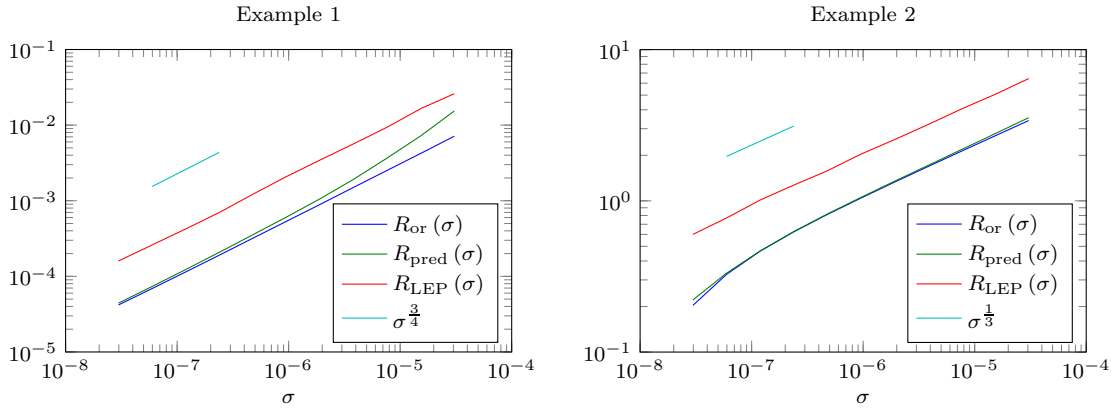

Figure 4: Simulation results for Showalter regularization.

prediction risk $R_{\mathrm{pred}}(\sigma)$, and the Lepskiĭ risk $R_{\mathrm{LEP}}(\sigma)$ for Tikhonov regularization. In Figure 5 we depict the fractions of the oracle risk with the different a-posteriori risks for various parameters ν and a to compare the average behavior of these parameter choice methods.

In conclusion we empirically find that both choices α_{pred} and α_{LEP} perform nearly as good as the oracle choice if σ is not too small. Comparing the performance of α_{pred} and α_{LEP} it seems

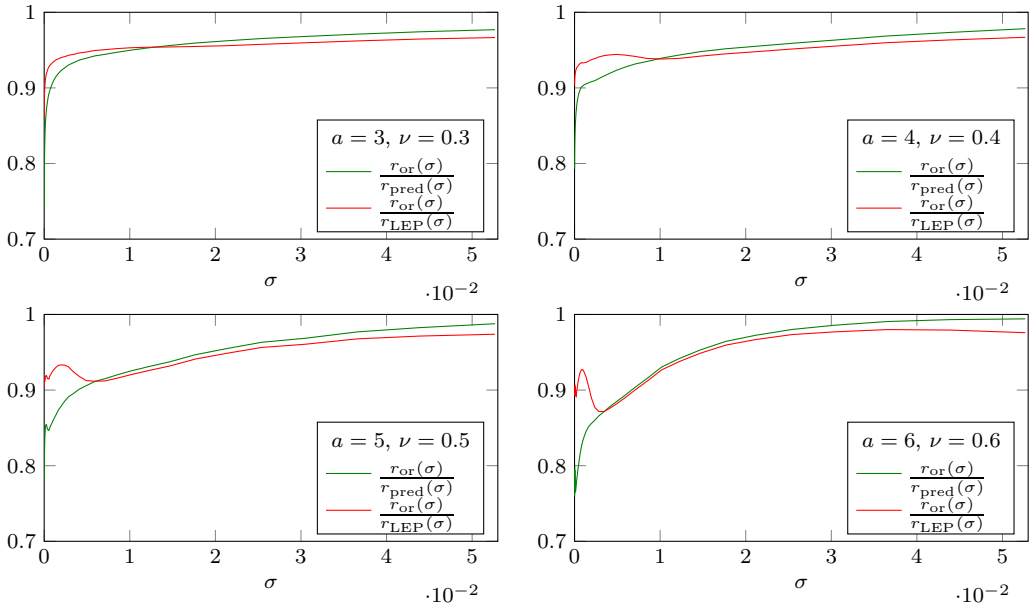

Figure 5: Efficiency simulations for Tikhonov regularization with different smoothness parameters a and ν.

that α_{LEP} behaves better for small values of σ, even though this has not been observed in Figures 2–4. Nevertheless, the computational effort for α_{pred} is significantly smaller: in our efficiency simulations around 90% of the computation time were spent for computing α_{LEP} in (22).

6 Conclusion and outlook

In this study we have analyzed a parameter choice method for filter-based regularization methods applied to linear statistical inverse problems. Therefore we have proven an oracle inequality, which generalizes the one from Chernousova and Golubev (2014) to general ordered filters satisfying weak assumptions (cf. Definition 1 and Assumption 3). From this oracle inequality we derived convergence rates of the investigated parameter choice, which are actually order optimal in a wide range of cases.

We point out that our techniques do not seem to be limited to the specific choice of α investigated here. Quite the contrary, we provide a general strategy to prove an oracle inequality, which might be used for other choices as well. If one would consider a different parameter choice rule, then an oracle inequality can be carried out the same way as in this study provided the following ingredients are available:

- general moment bounds for the prediction risk as in Corollary 4.1 (in our case based on the exponential bounds proven by Kneip (1994)),
- a deterministic optimality result as in Lemma 4.2,
- estimates for the behavior of the regularization algorithm as in Lemma 4.3.

Note that the latter two assumptions do not rely on the parameter choice rule at all but only on the interplay of the operator, the regularization method, and the source condition. Consequently, whenever the general moments can be bounded, the analysis can basically be carried over from standard results, see e.g. Bellec and Tsybakov (2016) for a useful deviation inequality for regularization methods with convex penalties.

Even though we have always assume that the noise level σ is known, the generalization to unknown σ is straightforward. For the choice of parameter α_{pred}, we simply replace σ by a proper estimator $\hat{\sigma}$ in (12). The only affected part in our argumentation is the Kneip's deviation bound in Theorem 2.1, which still holds if we further assume the smoothness of Tf (this is usually the case due to the blessing of ill-posedness), see (Kneip, 1994, Section 6) for further details.

Another possible generalization concerns the errors in (1). If ξ is such that the ξ_k's in (5) are independent sub-Gaussian errors, then there are two crucial parts of the proofs which have to be generalized: Theorem 2.1 and Lemma 4.4. In fact it turns out that both also hold for independent sub-Gaussian errors (see Kneip, 1994; Golubev, 2010), so the whole analytical methodology remains valid in such a case as well.

The general analytical strategy advocated in this paper, of course, has its own limitations, as the resulting oracle inequality might turn out to be inadequate or even trivial in certain cases (see Section 3.3.2 for instance). Future questions include generalizations to exponentially ill-posed problems, nonlinear problems and noise models with heavier tails.

Acknowledgements

FW wants to thank Yu. Golubev for pointing his attention to the inspiring paper (Chernousova and Golubev, 2014) during a visit in Göttingen. FW gratefully acknowledges financial support by the German Research Foundation DFG through subproject A07 of CRC 755, and HL acknowledges support through RTG 2088, subproject B2, and the National Nature Science Foundation of China (61571008, 61402495).

References

Abramovich, F. and Silverman, B. W. (1998). Wavelet decomposition approaches to statistical inverse problems. *Biometrika*, 85:115–129.

Bakushinskiĭ, A. B. (1984). Remarks on the choice of regularization parameter from quasioptimality and relation tests. *Zh. Vychisl. Mat. i Mat. Fiz.*, 24(8):1258–1259.

Bauer, F. and Hohage, T. (2005). A Lepskij-type stopping rule for regularized Newton methods. *Inverse Probl.*, 21(6):1975.

Bauer, F. and Lukas, M. A. (2011). Comparing parameter choice methods for regularization of ill-posed problems. *Math. Comput. Simulation*, 81(9):1795–1841.

Bellec, P. C. and Tsybakov, A. B. (2016). Bounds on the prediction error of penalized least squares estimators with convex penalty. *arXiv:1609.06675*.

Birgé, L. and Massart, P. (2001). Gaussian model selection. *J. Eur. Math. Soc. (JEMS)*, 3(3):203–268.

Birgé, L. and Massart, P. (2007). Minimal penalties for Gaussian model selection. *Probab. Theory Related Fields*, 138(1-2):33–73.

Bissantz, N., Hohage, T., Munk, A., and Ruymgaart, F. (2007). Convergence rates of general regularization methods for statistical inverse problems and applications. *SIAM J. Numer. Anal.*, 45(6):2610–2636.

Blanchard, G., Hoffmann, M., and Reiß, M. (2016). Optimal adaptation for early stopping in statistical inverse problems. *arXiv:1606.07702*.

Candès, E. J., Sing-Long, C. A., and Trzasko, J. D. (2013). Unbiased risk estimates for singular value thresholding and spectral estimators. *IEEE Trans. Signal Process.*, 61(19):4643–4657.

Cavalier, L. (2011). Inverse problems in statistics. In *Inverse problems and high-dimensional estimation*, volume 203 of *Lect. Notes Stat. Proc.*, pages 3–96. Springer, Heidelberg.

Cavalier, L., Golubev, G. K., Picard, D., and Tsybakov, A. B. (2002). Oracle inequalities for inverse problems. *Ann. Statist.*, 30(3):843–874. Dedicated to the memory of Lucien Le Cam.

Cavalier, L. and Golubev, Y. (2006). Risk hull method and regularization by projections of ill-posed inverse problems. *Ann. Statist.*, 34(4):1653–1677.

Cavalier, L., Golubev, Y., Lepski, O., and Tsybakov, A. (2003). Block thresholding and sharp adaptive estimation in severely ill-posed inverse problems. *Teor. Veroyatnost. i Primenen.*, 48(3):534–556.

Chernousova, E. and Golubev, Y. (2014). Spectral cut-off regularizations for ill-posed linear models. *Math. Methods Statist.*, 23(2):116–131.

Cohen, A., Hoffmann, M., and Reiß, M. (2004). Adaptive wavelet Galerkin methods for linear inverse problems. *SIAM J. Numer. Anal.*, 42(4):1479–1501.

Davies, A. R. and Anderssen, R. S. (1986). Improved estimates of statistical regularization parameters in Fourier differentiation and smoothing. *Numer. Math.*, 48(6):671–697.

Deledalle, C.-A., Vaiter, S., Fadili, J., and Peyré, G. (2014). Stein Unbiased GrAdient estimator of the Risk (SUGAR) for multiple parameter selection. *SIAM J. Imaging Sci.*, 7(4):2448–2487.

Dette, H., Munk, A., and Wagner, T. (1998). Estimating the variance in nonparametric regression-what is a reasonable choice? *J. R. Stat. Soc. Ser. B. Stat. Methodol.*, 60(4):751–764.

Donoho, D. L. (1995). Nonlinear solution of linear inverse problems by wavelet-vaguelette decomposition. *Appl. Comput. Harmon. Anal.*, 2(2):101–126.

Engl, H., Hanke, M., and Neubauer, A. (1996). *Regularization of Inverse Problems*. Springer.

Goldenshluger, A. (1999). On pointwise adaptive nonparametric deconvolution. *Bernoulli*, 5(5):907–925.

Golub, G. H., Heath, M., and Wahba, G. (1979). Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, 21(2):215–223.

Golubev, G. K. and Khas'minskiĭ, R. Z. (1999). A statistical approach to some inverse problems for partial differential equations. *Problemy Peredachi Informatsii*, 35(2):51–66.

Golubev, Y. (2010). On universal oracle inequalities related to high-dimensional linear models. *Ann. Statist.*, 38(5):2751–2780.

Hall, P., Kay, J. W., and Titterinton, D. M. (1990). Asymptotically optimal difference-based estimation of variance in nonparametric regression. *Biometrika*, 77:521–528.

Ingster, Y., Laurent, B., and Marteau, C. (2014). Signal detection for inverse problems in a multidimensional framework. *Math. Methods Statist.*, 23(4):279–305.

Ingster, Y. I., Sapatinas, T., and Suslina, I. A. (2012). Minimax signal detection in ill-posed inverse problems. *Ann. Statist.*, 40(3):1524–1549.

Johnstone, I. M., Kerkyacharian, G., Picard, D., and Raimondo, M. (2004). Wavelet deconvolution in a periodic setting. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 66(3):547–573.

Johnstone, I. M. and Silverman, B. W. (1991). Discretization effects in statistical inverse problems. *J. Complexity*, 7:1–34.

Kneip, A. (1994). Ordered linear smoothers. *Ann. Statist.*, 22(2):835–866.

Lepskiĭ, O. V. (1991). On a problem of adaptive estimation in Gaussian white noise. *Theory Probab. Appl.*, 35(3):454–466.

Li, K.-C. (1987). Asymptotic optimality for C_p, C_L, cross-validation and generalized cross-validation: discrete index set. *Ann. Statist.*, 15(3):958–975.

Lucka, F., Proksch, K., Brune, C., Bissantz, N., Burger, M., Dette, H., and Wübbeling, F. (2017). Risk estimators for choosing regularization parameters in ill-posed problems – properties and limitations. *arXiv:1701.04970*.

Luisier, F., Blu, T., and Unser, M. (2007). A new sure approach to image denoising: Interscale orthonormal wavelet thresholding. *IEEE Trans. Image Process.*, 16(3):593–606.

Lukas, M. A. (1993). Asymptotic optimality of generalized cross-validation for choosing the regularization parameter. *Numer. Math.*, 66(1):41–66.

Lukas, M. A. (1995). On the discrepancy principle and generalised maximum likelihood for regularisation. *Bull. Austral. Math. Soc.*, 52(3):399–424.

Mair, B. A. and Ruymgaart, F. H. (1996). Statistical inverse estimation in Hilbert scales. *SIAM J. Appl. Math.*, 56(5):1424–1444.

Mallows, C. L. (1973). Some comments on C_p. *Technometrics*, 15(4):661–675.

Mathé, P. (2006). The Lepskiĭ principle revisited. *Inverse Probl.*, 22(3):L11–L15.

Mathé, P. and Hofmann, B. (2008). How general are general source conditions? *Inverse Problems*, 24(1):015009.

Mathé, P. and Pereverzev, S. V. (2001). Optimal discretization of inverse problems in Hilbert scales. Regularization and self-regularization of projection methods. *SIAM J. Numer. Anal.*, 38(6):1999–2021.

Mathé, P. and Pereverzev, S. V. (2003). Geometry of linear ill-posed problems in variable Hilbert scales. *Inverse Probl.*, 19(3):789–803.

Mathé, P. and Pereverzev, S. V. (2006). Regularization of some linear ill-posed problems with discretized random noisy data. *Math. Comp.*, 75(256):1913–1929 (electronic).

Morozov, V. A. (1966). On the solution of functional equations by the method of regularization. *Soviet Math. Dokl.*, 7:414–417.

O'Sullivan, F. (1986). A statistical perspective on ill-posed inverse problems. *Statist. Sci.*, 1(4):502–527. With comments and a rejoinder by the author.

Pinsker, M. S. (1980). Optimal filtration of square-integrable signals in Gaussian noise. *Probl. Inf. Transm.*, 16(2):52–68. (Russian).

Rice, J. (1984). Bandwidth choice for nonparametric regression. *Ann. Statist.*, 12(4):1215–1230.

Rieder, A. (2005). Runge-Kutta integrators yield optimal regularization schemes. *Inverse Probl.*, 21(2):453–471.

Stein, C. M. (1981). Estimation of the mean of a multivariate normal distribution. *Ann. Statist.*, 9(6):1135–1151.

Tsybakov, A. (2000). On the best rate of adaptive estimation in some inverse problems. *C. R. Acad. Sci. Paris Sér. I Math.*, 330(9):835–840.

Vogel, C. R. (1986). Optimal choice of a truncation level for the truncated SVD solution of linear first kind integral equations when data are noisy. *SIAM J. Numer. Anal.*, 23(1):109–117.

Vogel, C. R. (2002). *Computational Methods for Inverse Problems*, volume 23 of *Frontiers in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA.

Wahba, G. (1977). Practical approximate solutions to linear operator equations when the data are noisy. *SIAM J. Numer. Anal.*, 14(4):651–667.

Wang, Y.-Q. and Morel, J.-M. (2013). SURE guided Gaussian mixture image denoising. *SIAM J. Imaging Sci.*, 6(2):999–1034.

Werner, F. (2015). On convergence rates for iteratively regularized Newton-type methods under a Lipschitz-type nonlinearity condition. *J. Inverse Ill-Posed P.*, 23(1):75–84.

Werner, F. and Hohage, T. (2012). Convergence rates in expectation for Tikhonov-type regularization of Inverse Problems with Poisson data. *Inverse Probl.*, 28(10):104004.